

FIG. 3




FIG. 8




FIG. 9

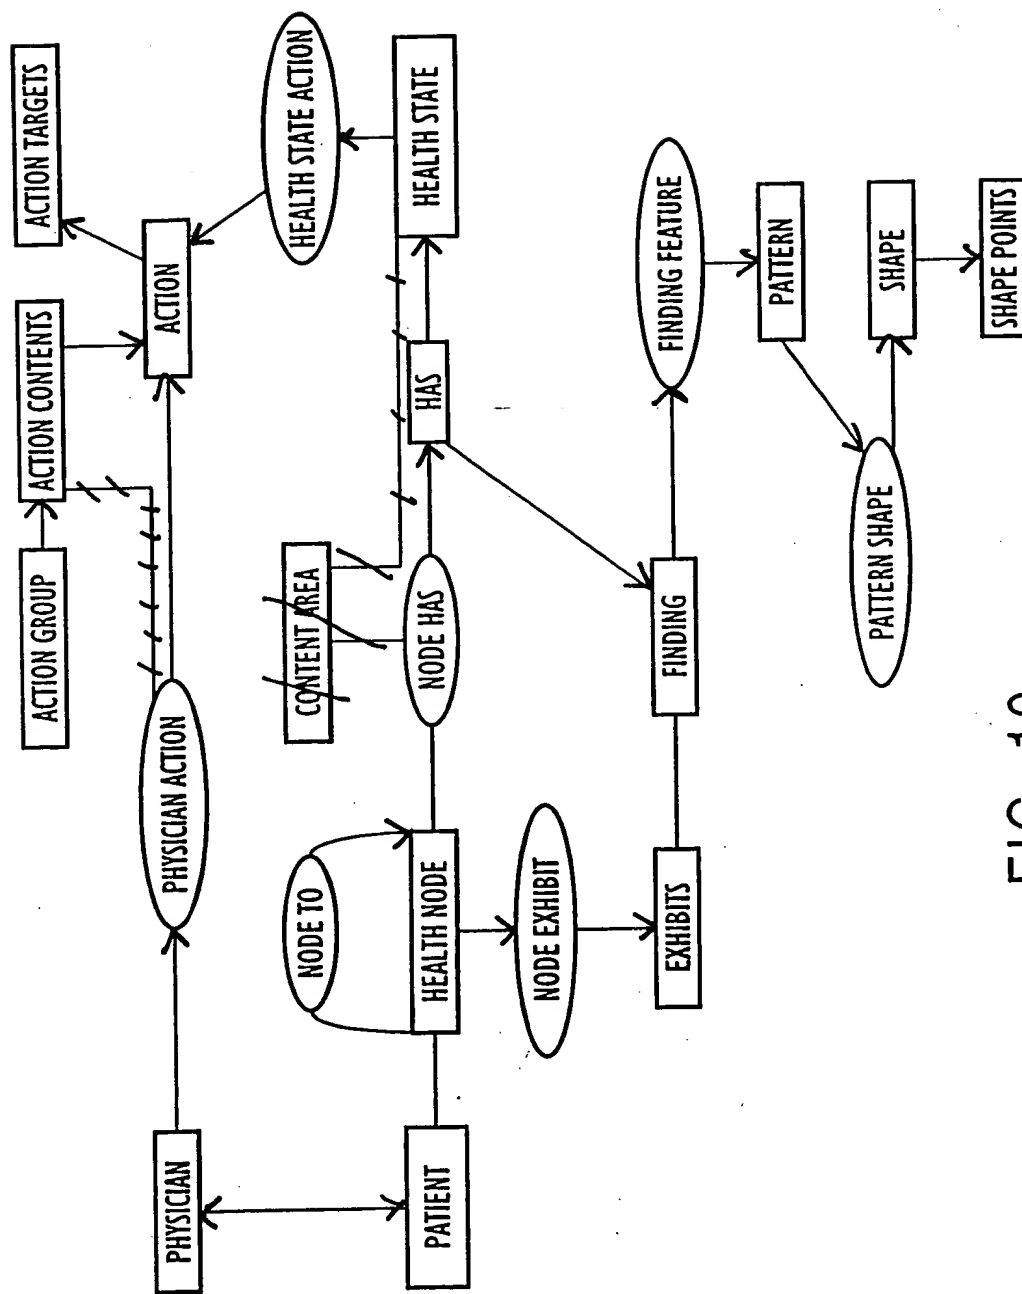


FIG. 10

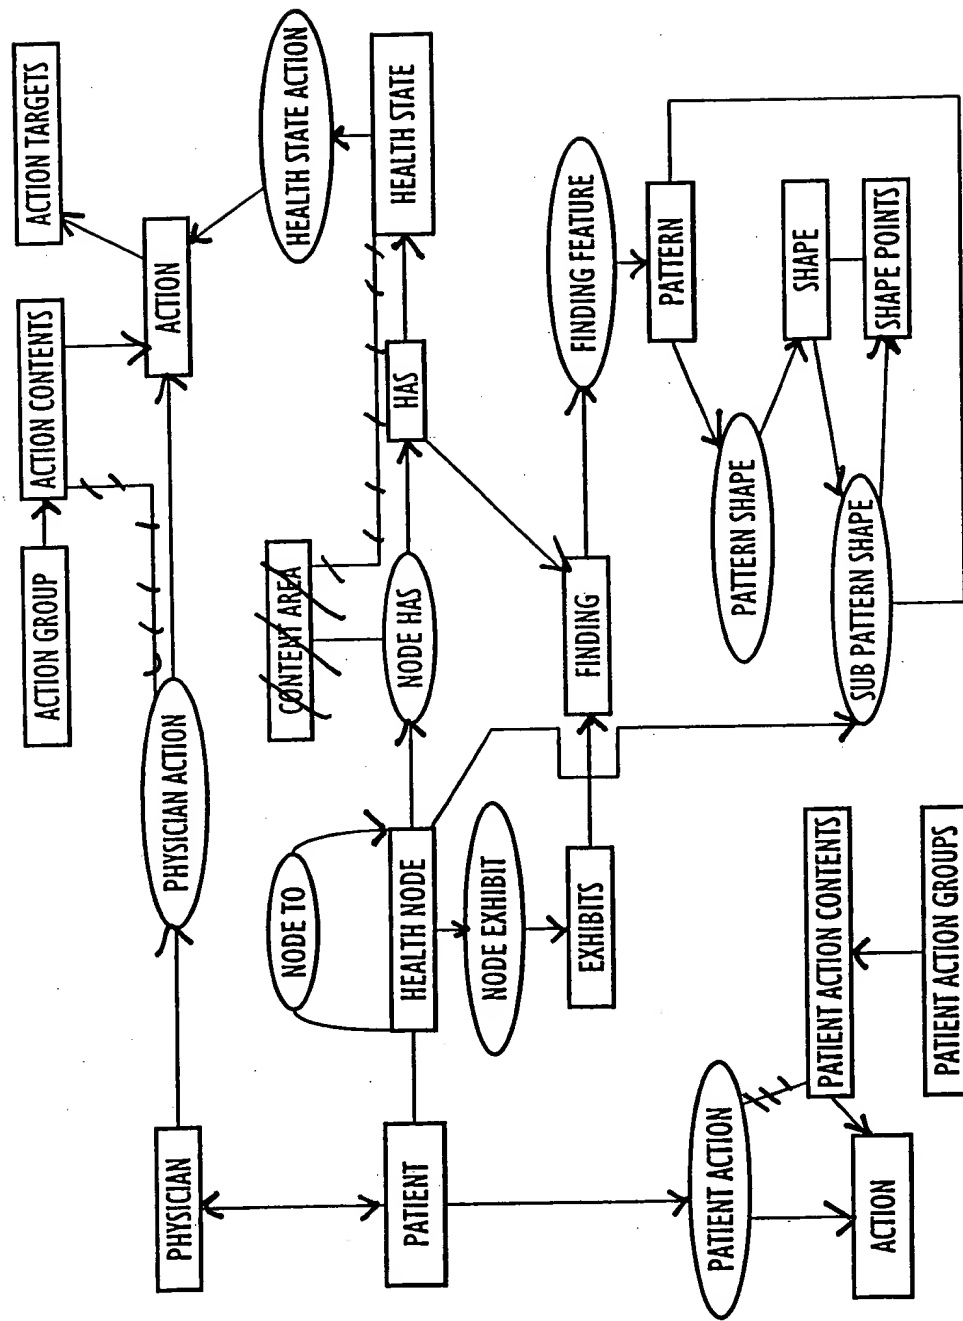


FIG. 11

START PROCESS

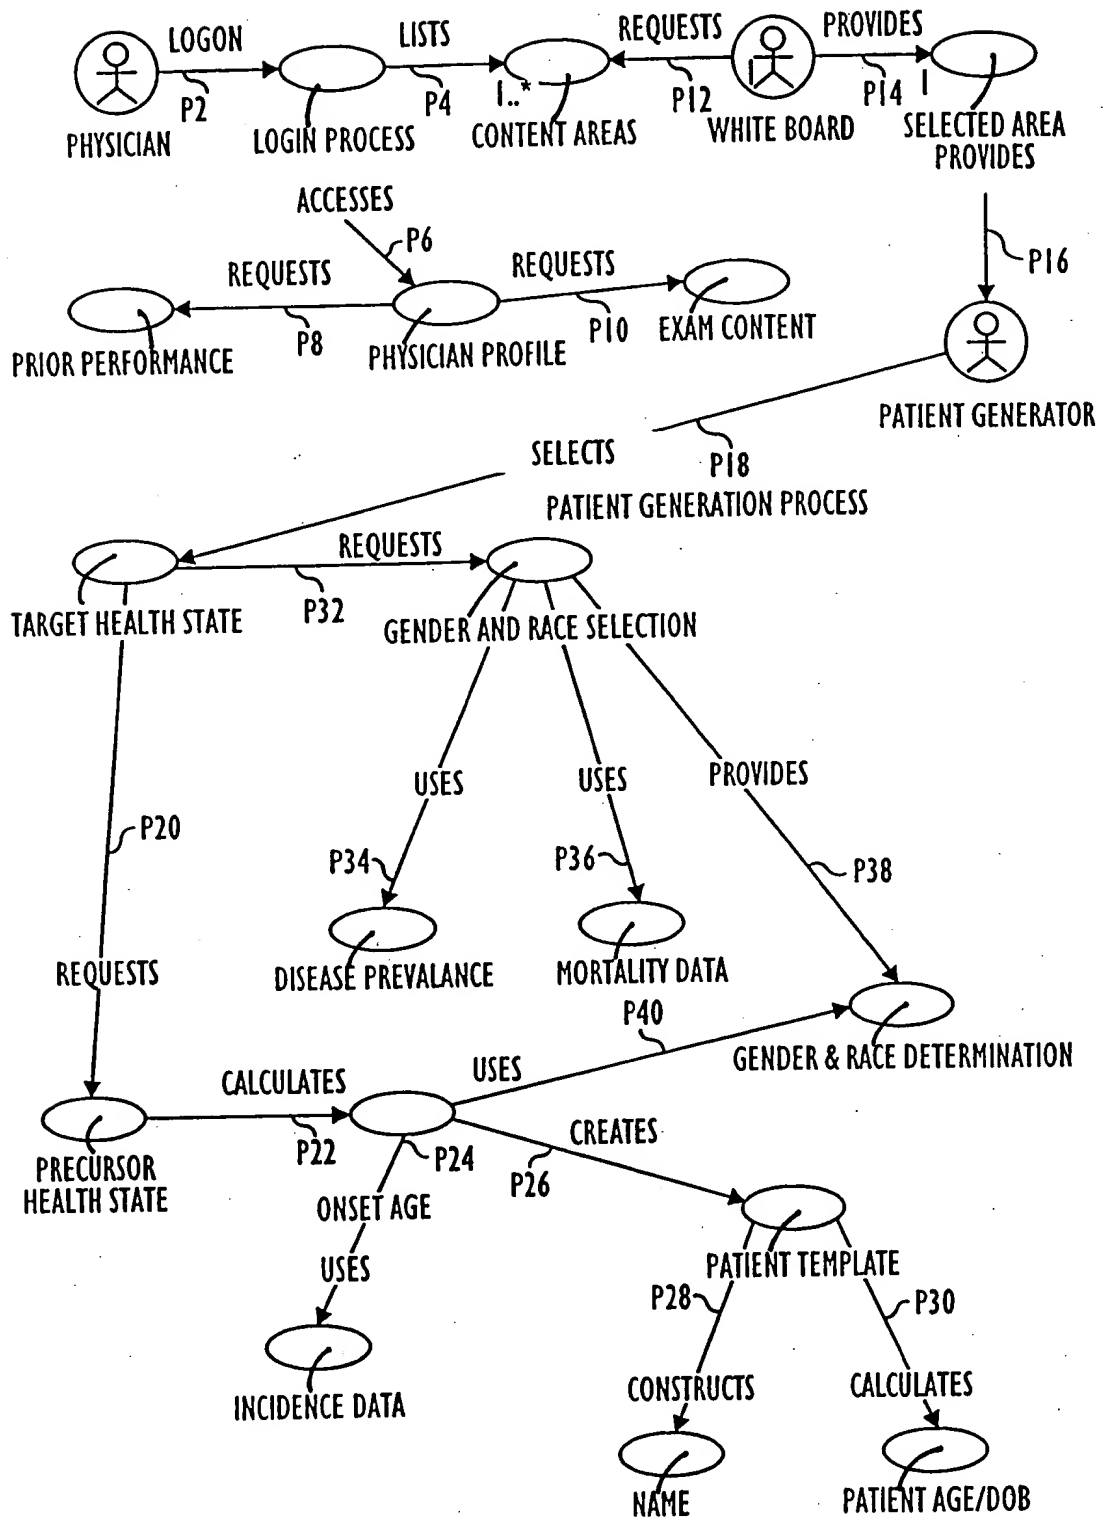


FIG. 13



FIG. 14




FIG. 15

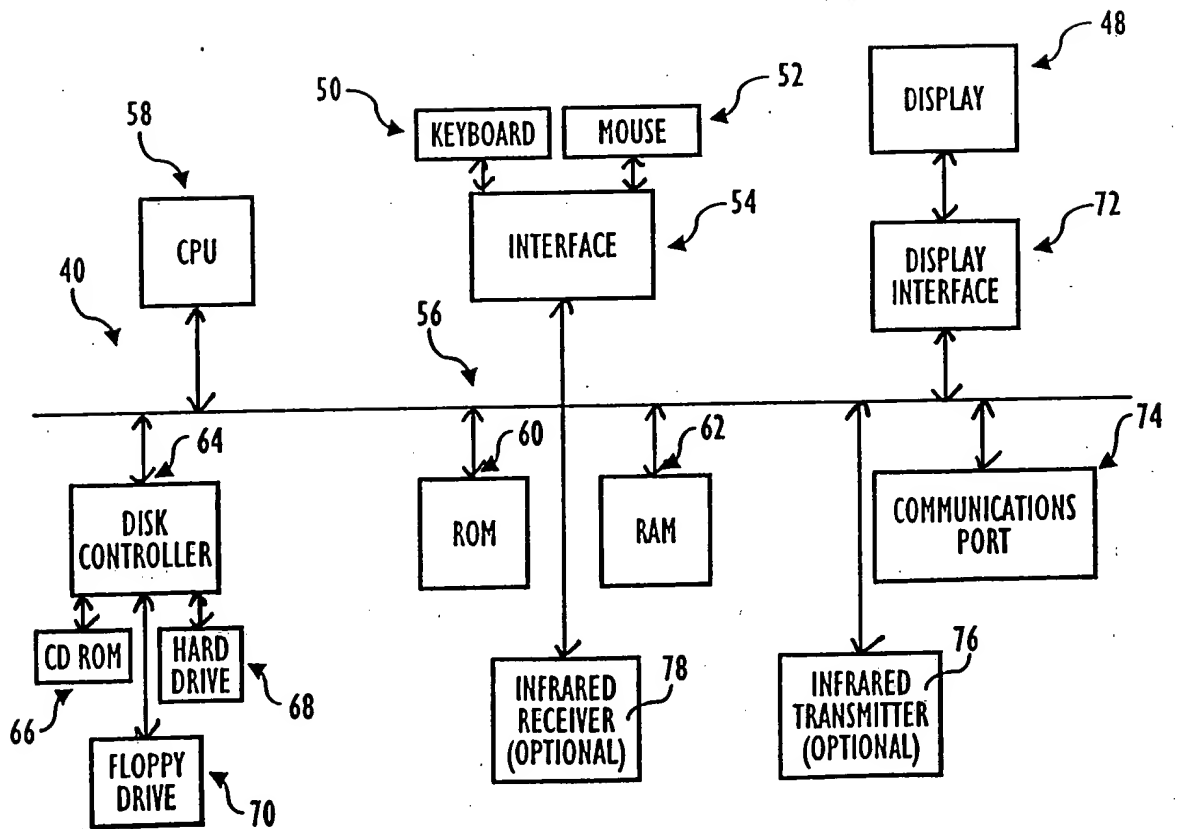


FIG. 16

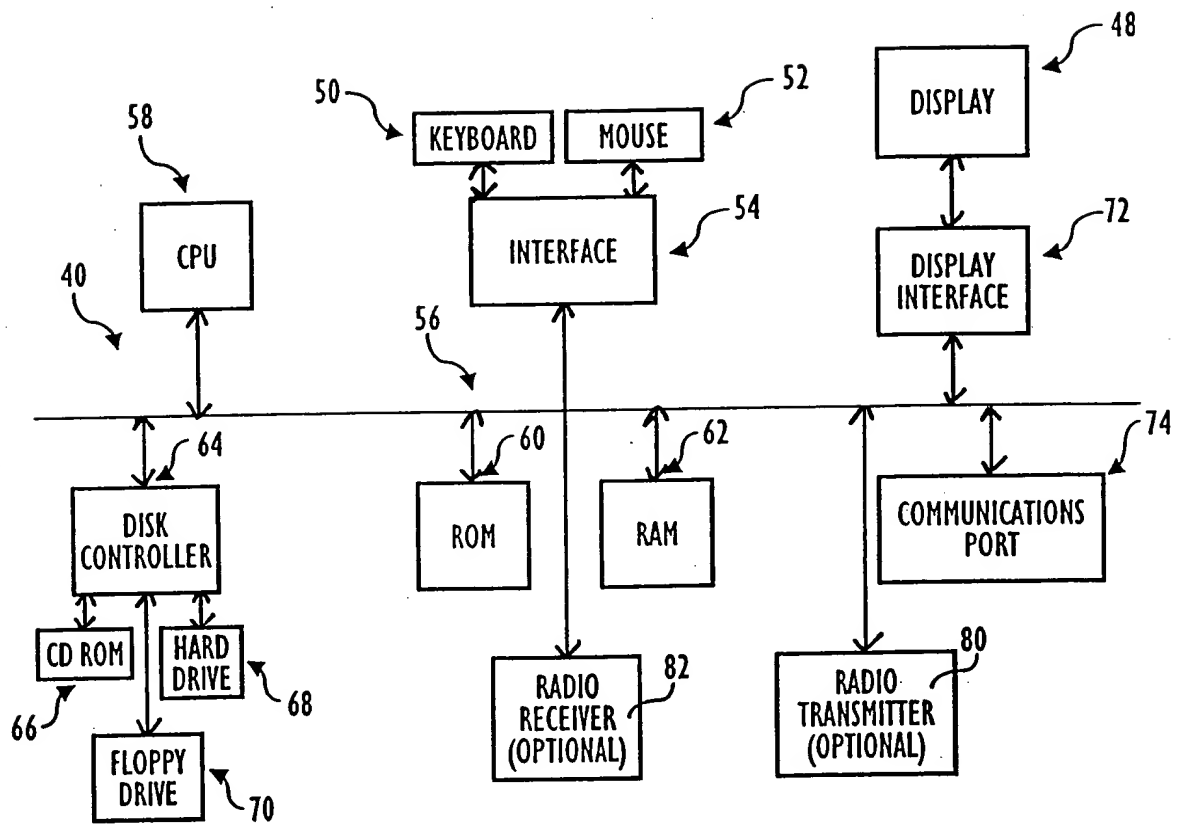
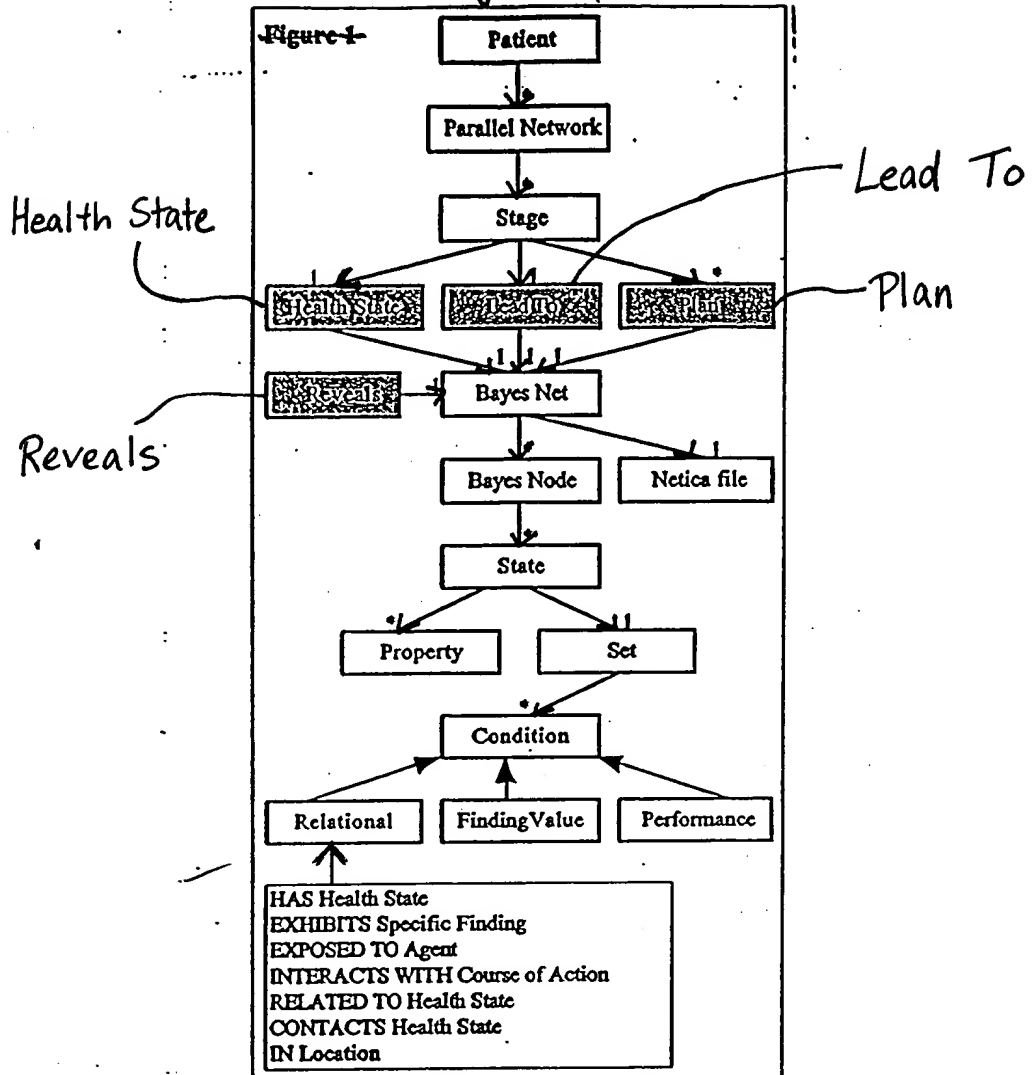


FIG. 17

Figure 19



Bayesian nets in this article, but the following fragments illustrate the use of Bayes nets and supporting structures in place of scripts.

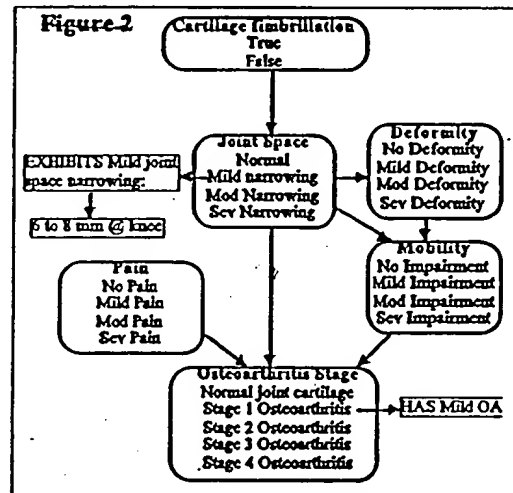
Figure 2 illustrates a simplified OA generating Bayes net. The network was built as a roughly physiologic model of the development of OA, assuming that cartilaginous deformities and destruction cause joint space narrowing, accompanied by sclerosis and subchondral cyst formation (not shown), and leading to gross deformities and loss of mobility. Pain is a variable feature, but probably must be present in a test case (otherwise, how would the doctor's attention be drawn to the joint?). The mild narrowing state of the joint space node is defined by a Set containing one Condition, EXHIBITS the Specific Finding, mild

BEST AVAILABLE COPY

joint space narrowing, which is itself defined as a joint space of 4 to 6 mm for the knee. The stage 1 state of the osteoarthritis stage node is similarly defined by a Set containing one Condition, HAS mild OA. The conditional probability tables for this node are arranged so that certain combinations of joint space narrowing and deformity define mild OA. Other combinations may define other stages, or be declared impossible (e.g. severe deformity without joint space narrowing might be impossible in this context).

Two interesting benefits of this approach are first, that we can use a single Bayes net to describe all five stages of OA, and second, through the logical magic of Bayes theorem, we can now invert a Bayes net built from a perspective of classifying stages of OA, and use it to generate descriptions of OA. We assert that the patient HAS any stage of OA, update probabilities throughout the network, and start stochastically assigning states to indeterminate nodes.

Figure 2.0



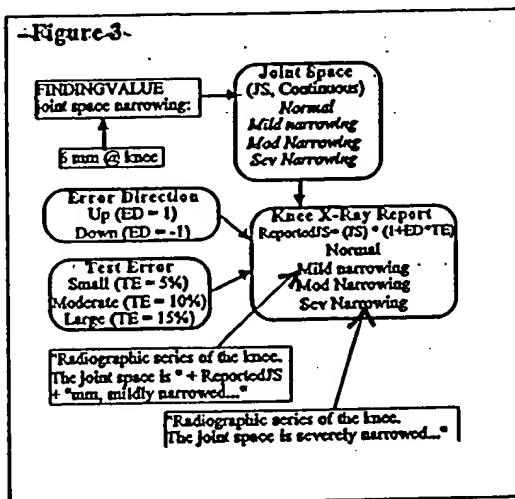
With each assignment, we write new information to the simulation, e.g. that the patient EXHIBITS mild joint space narrowing. We can test the Bayes net by experimenting with it from both perspectives, e.g. beginning with an assumption of cartilage damage to see what stage of OA results, or beginning with OA to see what other findings result.

Bayes nets supporting Leads to structures are conceptually very similar to Health State generating Bayes nets, with only two important differences. First, the Conditions usually describe risk factors for slower or more rapid progression, rather than features of a disease. For instance, an OA Lead To is likely to ask whether the patient HAS Obesity, or to assert that the patient is EXPOSED TO some remedy.

Second, the goal of the Lead to structure is to produce a rate of progression, which is not specified anywhere else in the simulation. (In contrast, the Health State generator has a goal of creating a description consistent with an asserted disease).

Figure 3 illustrates a Bayes net that could produce a report when an examinee requests a Revealing x-ray test. The only simulation data used in this report is the joint space, a value indirectly modified by the Bayes net in figure 2. Now we have a continuous Bayes node acquiring its value from a Condition that extracts the current joint space from the simulation. Note that there are no requests to determine whether any Specific Findings or Health States are present. Revealing queries should therefore be reusable across simulations. Also, the accuracy of the test can be built into the Bayes net representing that test. Another test, such as a magnetic resonance image, could have different size errors.

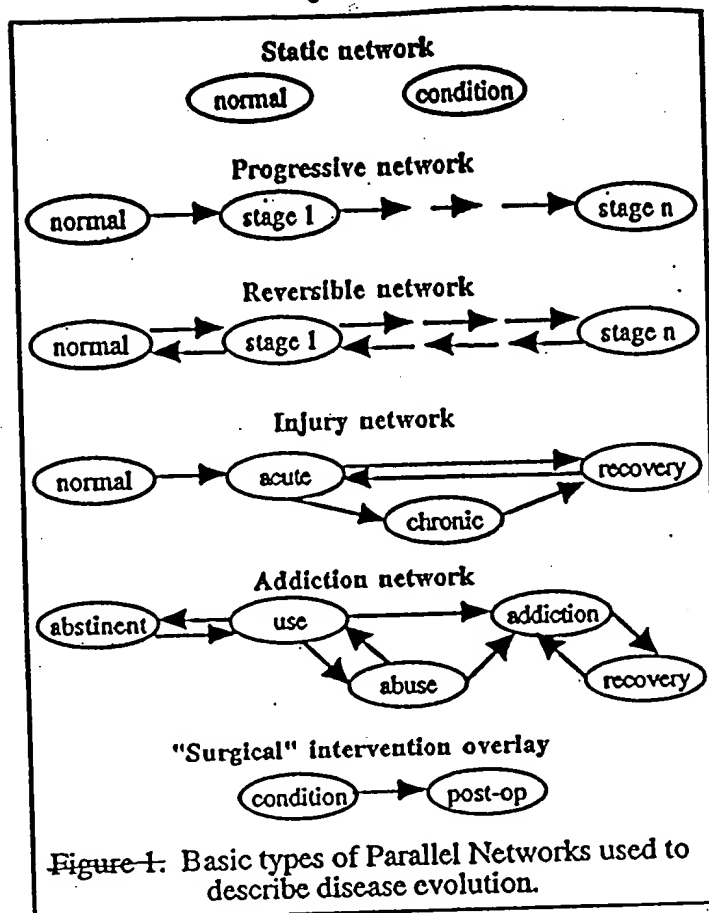
Figure 2)



Subjective queries are much more complex, but still possible to construct using the same approach. The primary complication is that subjective responses are uniquely elaborate in temporal detail, yielding statements such as, "the pain has been coming and going for weeks, but now it is worse than ever."

Management Plan critiques are similar to Reveals, except that most of the Conditions inspect prescriptions and queries made by the examinee, and the resulting report is a critique of a physicians' strategy. Our experience to date confirms the expectation that Bayes nets support inferences about actions and plans.⁶

Figure 22



disorders. In the injury network an acute insult evolves to either recovery or a chronic condition with a later recovery. Injury networks describe many infectious diseases and trauma. The addiction network illustrates that a person may abstain from, use, abuse, or become addicted to a substance. In the scheme shown here, a previously addicted person can only be addicted or recovering, but can not return to abstinence, use or abuse. The surgical intervention overlay illustrates that new states can be added to the above networks using irreversible therapies such as radiation or surgery. Domain experts adapted these networks to their needs by eliminating unwanted nodes and arcs, or replacing nodes with another network.

Domain experts began with a primary Parallel Network to sketch the diseases defining their domain, such as stages of diabetes mellitus. Parallel Networks of comorbid conditions were identified in most domains, typically including risk factors for

temperature must combine information about the underlying temperature and all antipyretic drugs administered in the last four hours.

The distinction between changing health states and perturbing findings is necessarily artificial (health states are just collections of findings), and the decision to model a particular process one way or another may often depend on testing goals, and subsequent decisions about how finely to model health states. In general, very fine distinctions between health states should result in more interventions that change health states, while coarsely defined health states may require more perturbations in findings.

A COA can modify the health state in which a patient exists at one point in time. When the candidate selects such a COA, the simulated patient may evolve to a new health state on the basis of patterns specified for health state evolution in the knowledge base. The knowledge for a particular health domain is stored as a parallel health state network. For example, the initially generated patient for a case of osteoarthritis will demonstrate some stage of osteoarthritis. However, other health states such as obesity might influence the progress of the patient's arthritis from mild to moderate and moderate to severe disease. In the parallel networks of health states representation, a newly-generated patient will display findings consistent with a health state in the primary domain (for example, osteoarthritis) and in the parallel health states (eg, obesity) which influence the primary health state's progress. As shown in the following figure, osteoarthritis can progress over time from the normal state to mild, moderate or severe osteoarthritis.

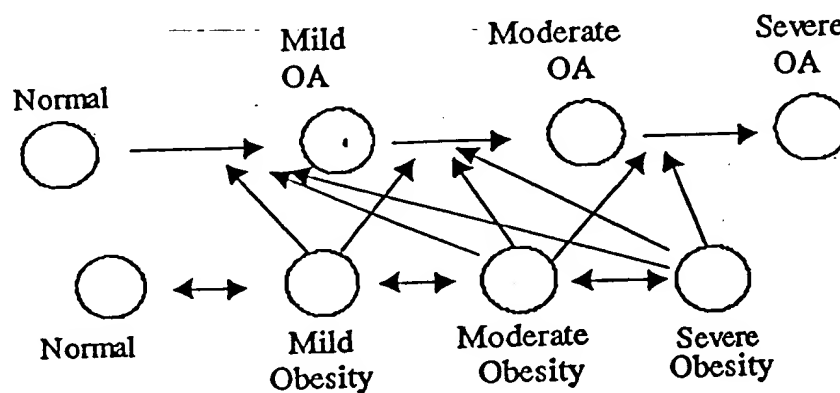


Figure 23


For this particular illness, progress occurs in one direction only; osteoarthritis doesn't regress once developed, but can stabilize at a particular degree of severity. Obesity represents a parallel health state which can influence the progression of osteoarthritis. Mild, moderate, and severe obesity can influence this progress at different rates: the model permits representation of greater impact for more severe obesity states. Notice also that obesity can regress (severe obesity can revert to moderate obesity, etc.) Similarly, other parallel health states might exist which could modify progression of osteoarthritis. For example, the patient who has osteoarthritis will frequently utilize non-steroidal anti-inflammatory drugs (NSAID's) for treatment. These agents can improve the symptoms of osteoarthritis, but also impact on the parallel state of peptic ulcer disease, ie treatment with NSAID's can induce an ulcer, which will then evolve in parallel with the course and treatment of osteoarthritis. Initial experience with this representation indicates that these modifier-relationships are not well-defined in the medical literature and constitute a research area for further development.

The simulation system's fidelity depends upon access to a rich representation of health state-specific knowledge. This knowledge consists of Findings obtained from physician "knowledge-donors" working from templates provided by the Assessment Technologies, Inc. development team. The template includes a NAME for the health state and an associated SNOMED code. The template also includes specific descriptions of the *Findings*, and *Patterns* for these *Findings*. The patterns are stored as distributions; these distributions are obtained from the medical literature where available, and from physician expert opinion where such published data don't exist.

The development team doesn't expect the knowledge groups to provide these distributions but rather to indicate the relationships between health states, how parallel states influence each other qualitatively (eg increase, decrease, stay the same), and possible sources of information about the relevant probabilities.

The knowledge model has evolved over the past six months to include extensive use of belief networks (also called Bayesian networks). Belief networks provide a graphical process for describing the relationships between entities in a health state. For example, some set of characteristics (family history, age, gender, racial origin/ethnicity, body weight) influence the development of impaired glucose tolerance.

Figure 2.4



This "graph" illustrates these relationships: Family History, Gender, Race, Weight, and Age all influence the development of impaired glucose tolerance. The raw pictorial doesn't say how they influence IGT, but rather that they "influence" the development of IGT. In the background, we incorporate probabilistic information which describes these relationships quantitatively, but would expect the knowledge development group to provide only semi-quantitative guidance (eg a person whose mother has diabetes has twice the likelihood of developing IGT compared to an individual who has no such family history.) We intend to fill in the more specific quantitative probabilities on the basis of data in the literature where available; if such information don't exist, we will have to rely on expert opinion.

How The Knowledge Development Process Will Work

What will we need from the knowledge team in order to generate the information required in our system? The team should proceed in a step-wise fashion to address the following issues:

